Red Hat Hybrid Cloud Console 2023

Planning for Red Hat Hybrid Cloud Console

Planning Red Hat Hybrid Cloud Console Administration

Last Updated: 2023-05-02

Red Hat Hybrid Cloud Console 2023 Planning for Red Hat Hybrid Cloud Console

Planning Red Hat Hybrid Cloud Console Administration

Red Hat Customer Content Services

Legal Notice

Abstract

This guide provides an overview of Red Hat Hybrid Cloud Console features and can help Red Hat Enterprise Linux administrators with the planning and usage of the many services included in the Hybrid Cloud Console. Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see our CTO Chris Wright's message.

Table of Contents

CHAPTER 1. PLANNING USER ACCESS

The User Access feature is an implementation of role-based access control (RBAC) that controls user access to various services hosted on the Red Hat Hybrid Cloud Console. You configure the User Access feature to grant user access to services hosted on Hybrid Cloud Console.

1.1. WHO CAN USE USER ACCESS

To initially view and manage User Access on Red Hat Hybrid Cloud Console, you must be an Organization Administrator. This is because User Access requires user management capabilities that are designated from the Red Hat Customer Portal at Customer Portal. Those capabilities belong solely to the Organization Administrator.

The **User Access administrator** role is a special role that the Organization Administrator can assign. This role allows users who are not Organization Administrator users to manage User Access on Red Hat Hybrid Cloud Console.

1.2. ADDITIVE ACCESS

User access on Red Hat Hybrid Cloud Console uses an additive model, which means that there are no **deny** roles. In other words, actions are only permitted. To control access, assign the appropriate roles with the desired permissions to groups, then add users to those groups. The access permitted to any individual user is a sum of all roles assigned to all groups to which that user belongs.

1.3. THE USER ACCESS GROUPS, ROLES, AND PERMISSIONS

User Access uses the following categories to determine the level of user access that an Organization Administrator can grant to the supported Red Hat Hybrid Cloud Console services. The access provided to any authorized user depends on the group that the user belongs to and the roles assigned to that group.

- **Group**: A collection of users belonging to an account which provides the mapping of roles to users. An Organization Administrator can use groups to assign one or more roles to a group and to include one or more users in a group. You can create a group with no roles and no users.
- **Roles**: A set of permissions that provide access to a given service, such as Insights. The permissions to perform certain operations are assigned to specific roles. Roles are assigned to groups. For example, you might have a **read** role and a **write** role for a service. Adding both roles to a group grants all members of that group read and write permissions to that service.
- **Permissions**: A discrete action that can be requested of a service. Permissions are assigned to roles.

An Organization Administrator adds or deletes roles and users to groups. The group can be a new group created by an Organization Administrator or the group can be an existing group. By creating a group that has one or more specific roles and then adding users to that group, you control how that group and its members interact with the Red Hat Hybrid Cloud Console services.

When you add users to a group, they become members of that group. A group member inherits the roles of all other groups they belong to. The user interface lists users in the **Members** tab.

CHAPTER 2. CONFIGURING SOURCES FOR RED HAT SERVICES

A data source is a service, application, or provider that supplies data to a Red Hat Hybrid Cloud Console application or service. Sources comprise cloud sources and Red Hat sources. The services and applications on the Red Hat Hybrid Cloud Console use sources to connect with public cloud providers and other services or tools to collect information for the service or application.

You add and manage sources in the **Sources** application located within the **Settings** bundle. To access **Settings**, click the gear icon in the masthead on the Red Hat Hybrid Cloud Console.

The **Sources** menu uses a wizard to help you add cloud sources and Red Hat sources. For cloud sources, you can associate the provider with Red Hat applications, such as Cost Management and the RHEL management bundle. For Red Hat sources, you can add Red Hat OpenShift Container Platform. Adding applications is optional for cloud sources but is required for Red Hat sources.

2.1. ADDING CLOUD SOURCES

The **Add a cloud source** wizard steps you through creating a source. You can add Amazon Web Services (AWS), Google Cloud, and Microsoft Azure. The wizard provides detailed information for each public cloud provider.

Amazon Web Services

The workflow for adding AWS as a cloud source includes the following high-level steps:

1. Selecting the source type
2. Naming the source
3. Selecting the configuration
4. Selecting applications
5. Configuring cost and usage reporting (for cost management)
6. Identifying tags, aliases, and organizational units (for cost management)
7. Enabling account access
8. Reviewing details
9. Adding the source

You have two choices for the configuration mode:

- **Account authorization** (recommended)
- **Manual configuration**

If you select **Account authorization**, you provide your AWS account credentials (Access key ID and Secret key ID) and Red Hat configures and manages your source for you. This option automatically selects the **Cost Management** and the **RHEL management** bundle applications. You can deselect these applications.

If you select **Manual configuration**, you choose **Cost Management**, **RHEL management** bundle, or **No application**.

The **Cost Management** application allows you to perform financially related tasks, such as:

- Visualizing, understanding, and analyzing the use of resources and costs
- Forecasting your future consumption and comparing them with budgets
- Optimizing resources and consumption
- Identifying patterns of usage for further analysis
- Integrating with third-party tools that can benefit from cost and resourcing data

The **RHEL management** bundle includes the following items:

- Red Hat gold images
- High precision subscription watch data
- Autoregistration

The **Cost Management** and the **RHEL management** bundle applications require you to enable account access. You accomplish this by creating an IAM policy, an IAM role, and entering your Amazon Resource Name (ARN). An ARN is a generic name for an Amazon resource and has a common format depending on the service involved. In this case, it is the identity and access management (IAM) service and Role resource-type.

If you select **No application**, you choose which credentials to supply:

- AWS Secret key
- Cost Management ARN
- Subscription Watch ARN

Google Cloud

The workflow for adding Google Cloud as a cloud source includes the following high-level steps:

1. Selecting the source type
2. Naming the source
3. Selecting applications
4. Adding a project
5. Enabling account access
6. Creating a dataset
7. Setting up billing export information
8. Reviewing details
9. Adding the source

Cost Management is the only application choice. You must create an IAM role and assign access.

The **Cost Management** application allows you to perform financially related tasks, such as:

- Visualizing, understanding, and analyzing the use of resources and costs
- Forecasting your future consumption and comparing them with budgets
- Optimizing resources and consumption
- Identifying patterns of usage for further analysis
- Integrating with third-party tools that can benefit from cost and resourcing data

If you select **No application**, you provide the Project ID and the Service Account JSON as credentials.

Microsoft Azure

The workflow for adding Microsoft Azure as a cloud source includes the following high-level steps:

1. Selecting the source type
2. Naming the source
3. Selecting applications
4. Creating a resource group and a storage account (for cost management)
5. Entering a Subscription ID (for cost management)
6. Creating roles (for cost management)
7. Setting up daily exports (for cost management)
8. Providing credentials
9. Reviewing details
10. Adding the source

Application choices include **Cost Management**, **RHEL management** bundle, or **No application**.

The **Cost Management** application allows you to perform financially related tasks, such as:

- Visualizing, understanding, and analyzing the use of resources and costs
- Forecasting your future consumption and comparing them with budgets
- Optimizing resources and consumption
- Identifying patterns of usage for further analysis
- Integrating with third-party tools that can benefit from cost and resourcing data

The **RHEL management** bundle includes the following items:

- Red Hat gold images

- Autoregistration

You create a dedicated resource group and a storage account in the Azure Portal so you can collect cost data and metrics for cost management. You then use your subscription ID to create a Cost Management Storage Account Contributor role in the Cloud Shell. Using the subscription ID to run a second command in the Cloud Shell gives you the tenant (directory) ID, client (application) ID, and client secret necessary to complete setting up that role.

NOTE

Configure dedicated credentials to grant cost management read-only access to Azure cost data.

2.2. ADDING RED HAT SOURCES

The workflow for adding Red Hat OpenShift Container Platform as a cloud source includes the following high-level steps:

1. Selecting the source type and application
2. Naming the source
3. Installing and configuring the operator
4. Reviewing the details
5. Adding the source

Cost Management is the only application choice. The **Cost Management** application allows you to perform financially related tasks, such as:

- Visualizing, understanding, and analyzing the use of resources and costs
- Forecasting your future consumption and comparing them with budgets
- Optimizing resources and consumption
- Identifying patterns of usage for further analysis
- Integrating with third-party tools that can benefit from cost and resourcing data

NOTE

For Red Hat OpenShift Container Platform 4.6 and later, install the **costmanagement-metrics-operator** from the OpenShift Container Platform web console. For more information, see Adding an OpenShift Container Platform source to cost management .

2.3. SOURCES REFERENCE MATERIAL

Getting started with cost management

Getting Started with the Subscriptions Service

Adding sources for public cloud metering

CHAPTER 3. MANAGING SYSTEM TAGS AND GROUPS

Red Hat Hybrid Cloud Console enables administrators to filter groups of systems in inventory and in individual services using group tags. Groups are identified by the method of system data ingestion to Hybrid Cloud Console. Hybrid Cloud Console enables filtering groups of systems by those running SAP workloads, by Satellite host group, by Microsoft SQL Server workload, and by custom tags defined by system administrators who have root access to configure the Insights client on the system.

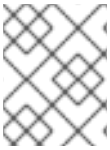

NOTE

As of Spring 2022, inventory, advisor, compliance, vulnerability, patch, drift, and policies enable filtering by groups and tags. Other services will follow.

IMPORTANT

Unlike the other services that enable tagging, the compliance service sets tags within lists of systems in the compliance service UI. For more information, see Group and tag filters in the compliance service.

Use the global, **Filter results** box to filter by SAP workloads, Satellite host groups, MS SQL Server workloads, or by custom tags added to the Insights client configuration file.

Prerequisites

The following prerequisites and conditions must be met to use the tagging features in Red Hat Hybrid Cloud Console:

- The Red Hat Insights client is installed and registered on each system.
- You must have root permissions, or their equivalent, to create custom tags or to change the **/etc/insights-client/tags.yaml** file.

3.1. SAP WORKLOADS

As Linux becomes the mandatory operating system for SAP ERP workloads in 2025, Red Hat Enterprise Linux and Red Hat Hybrid Cloud Console are working to make Hybrid Cloud Console the management tool of choice for SAP administrators.

As part of this ongoing effort, Hybrid Cloud Console automatically tags systems running SAP workloads and by SAP ID (SID), without any customization needed by administrators. Users can easily filter those workloads throughout the Hybrid Cloud Console application by using the global **Filter by tags** drop-down menu.

3.2. SATELLITE HOST GROUPS

Satellite host groups are configured in Satellite and recognized automatically by Hybrid Cloud Console.

3.3. MICROSOFT SQL SERVER WORKLOADS

Red Hat Hybrid Cloud Console users can use the global **Filter by tags** feature to select groups of systems running Microsoft SQL Server workloads.

In May of 2019, the Red Hat Insights team introduced a new set of Hybrid Cloud Console recommendations for Microsoft SQL Server running on Red Hat Enterprise Linux (RHEL). These rules

alert administrators to operating system level configurations that do not conform to the documented recommendations from Microsoft and Red Hat.

A limitation of these rules was that they primarily analyzed the operating system, and not the database itself. The latest release of Hybrid Cloud Console and RHEL 8.5 introduces Microsoft SQL Assessment API. The SQL Assessment API provides a mechanism to evaluate the database configuration of MS SQL Server for best practices. The API is delivered with a rule set that contains best practice rules that have been suggested by the Microsoft SQL Server Team. While this rule set is enhanced with the release of new versions, the API is built with the intent to give a highly customizable and extensible solution. This enables users to tune the default rules and to create their own.

The SQL Assessment API is supported by PowerShell for Linux (available from Microsoft), and Microsoft has developed a PowerShell script that can be used to call the API and store its results as a JSON formatted file. With RHEL 8.5, the Insights client now uploads this JSON file and presents the results in an easy-to-understand format in the Hybrid Cloud Console UI.

For more information about SQL Server assessment in Hybrid Cloud Console, see SQL Server database best practices now available through Red Hat Insights.

3.3.1. Setting up SQL Server assessments

To configure the Microsoft SQL Assessment API to provide information to Red Hat Insights, the database administrator needs to take the following steps.

Procedure

1. In the database you wish to assess, create a login for SQL Server assessments using SQL Authentication. The following Transact-SQL creates a login. Replace **<*PASSWORD*>** with a strong password:

```
USE [master]
GO
CREATE LOGIN [assessmentLogin] with PASSWORD= N'<*PASSWORD*>'
ALTER SERVER ROLE [sysadmin] ADD MEMBER [assessmentLogin]
GO
```

2. Store the credentials for login on the system as follows, again replacing <*PASSWORD*> with the password you used in step 1.

```
# echo "assessmentLogin" > /var/opt/mssql/secrets/assessment
# echo "<*PASSWORD*>" >> /var/opt/mssql/secrets/assessment
```

3. Secure the credentials used by the assessment tool by ensuring that only the mssql user can access the credentials.

```
# chmod 0600 /var/opt/mssql/secrets/assessment
# chown mssql:mssql /var/opt/mssql/secrets/assessment
```

4. Download PowerShell from the microsoft-tools repository. This is the same repository you configured when you installed the **mssql-tools** and **mssqlodbc17** packages as part of SQL Server installation.

```
# yum -y  install powershell
```

5. Install the SQLServer module for PowerShell. This module includes the assessment API.

```
# su mssql -c "/usr/bin/pwsh -Command Install-Module SqlServer"
```

6. Download the runassessment script from the Microsoft examples GitHub repository. Ensure it is owned and executable by mssql.

```
# /bin/curl -LJ0 -o /opt/mssql/bin/runassessment.ps1
https://raw.githubusercontent.com/microsoft/sql-server-samples/master/samples/manage/sql-assessment-api/RHEL/runassessment.ps1
# chown mssql:mssql /opt/mssql/bin/runassessment.ps1
# chmod 0700 /opt/mssql/bin/runassessment.ps1
```

7. Create the directory that will store the log file used by Red Hat Insights. Again, make sure it is owned and executable by mssql.

```
# mkdir /var/opt/mssql/log/assessments/
# chown mssql:mssql /var/opt/mssql/log/assessments/
# chmod 0700 /var/opt/mssql/log/assessments/
```

8. You can now create your first assessment, but be sure to do so as the user mssql so that subsequent assessments can be run automatically via cron or systemd more securely as the mssql user.

```
# su mssql -c "pwsh -File /opt/mssql/bin/runassessment.ps1"
```

9. Hybrid Cloud Console will automatically include the assessment next time it runs, or you can initiate Insights client by running this command:

```
# insights-client
```

3.3.1.1. Setting up the SQL Assessment on a timer

Because SQL Server Assessments can take 10 minutes or more to complete, it may or may not make sense for you to run the assessment process automatically every day. If you would like to run them automatically, the Red Hat SQL Server community has created **systemd** service and timer files to use with the assessment tool.

Procedure

1. Download the following files from Red Hat public SQL Server Community of Practice GitHub site.
   - **mssql-runassessment.service**
   - **mssql-runassessment.timer**
2. Install both files in the directory **/etc/systemd/system/**.

```
# cp mssql-runassessment.service /etc/systemd/system/
# cp mssql-runassessment.timer /etc/systemd/system/
# chmod 644 /etc/systemd/system/
```

3. Enable the timer.

```
# systemctl enable --now mssql-runassessment.timer
```

3.4. CUSTOM SYSTEM TAGGING

By applying custom grouping and tagging to your systems, you can add contextual markers to individual systems, filter by those tags in the Hybrid Cloud Console application, and more easily focus on related systems. This functionality can be especially valuable when deploying Hybrid Cloud Console at scale, with many hundreds or thousands of systems under management.

In addition to the ability to add custom tags to several Hybrid Cloud Console services, you can add predefined tags. The advisor service can use those tags to create targeted recommendations for your systems that might require more attention, such as those systems that require a higher level of security.

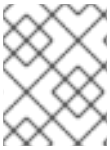

NOTE

To create custom and predefined tags, you must have root permissions, or their equivalent, to add to, or change the **/etc/insights-client/tags.yaml** file.

3.4.1. Tag structure

Tags use a **namespace/key=value** paired structure.

- **Namespace.** The namespace is the name of the ingestion point, *insights-client*, and cannot be changed. The **tags.yaml** file is abstracted from the namespace, which is injected by the Insights client before upload.
- **Key.** You can create the key or choose a predefined key from the system. Key names can use a mix of capitalization, letters, numbers, symbols and whitespace.
- **Value.** Define your own descriptive string value. You can use a mix of capitalization, letters, numbers, symbols and whitespace.

NOTE

The advisor service includes Red Hat-supported predefined tags.

3.4.2. Creating a tags.yaml file and adding a custom group

To create and add tags to **/etc/insights-client/tags.yaml** use **insights-client --group=<name-you-choose>**. This command performs the following actions:

- Creates the **etc/insights-client/tags.yaml** file
- Adds the **group=** key and **<name-you-choose>** value to **tags.yaml**
- Uploads a fresh archive from the system to the Hybrid Cloud Console application so the new tag is immediately visible, along with your latest results

After creating the initial **group** tag, add additional tags as needed by editing the **/etc/insights-client/tags.yaml** file.

The following procedure shows how to create the **/etc/insights-client/tags.yaml** file and the initial group, then how to verify that the tag exists in the Hybrid Cloud Console inventory.

Creating a new group

1. Run the following command as root, adding your custom group name after **--group=**:

```
[root@server ~]# insights-client --group=<name-you-choose>
```

Example of tags.yaml format

The following example of a **tags.yaml** file shows an example of file format and additional tags added for the new group:

```
# tags
---
group: eastern-sap
name: Jane Example
contact: jexample@corporate.com
Zone: eastern time zone
Location:
- gray_rack
- basement
Application: SAP
```

Verifying that your custom group was created

1. Navigate to Red Hat Enterprise Linux > Inventory and log in if necessary.
2. Click the **Filter results** dropdown menu.
3. Scroll through the list or use the search function to locate the tag.
4. Click the tag to filter by it.
5. Verify that your system is among the results on the advisor systems list.
6. Verifying that the system is tagged
7. Navigate to Red Hat Enterprise Linux > Inventory and log in if necessary.
8. Activate the **Name** filter and begin typing the system name until you see your system, then select it.
9. Verify that, next to the system name, the tag symbol is darkened and shows a number representing the correct number of tags applied.

3.4.3. Editing tags.yaml to add or change tags

After creating the group filter, edit the contents of **/etc/insights-client/tags.yaml** as needed to add or modify tags.

Procedure

1. Using the command line, open the tag configuration file for editing.

```
[root@server ~]# vi /etc/insights-client/tags.yaml
```

2. Edit content or add additional values as needed. The following example shows how you can organize **tags.yaml** when adding multiple tags to a system.

```
# tags
---
group: eastern-sap
location: Boston
description:
- RHEL8
- SAP
key 4: value
```

NOTE

Add as many key=value pairs as you need. Use a mix of capitalization, letters, numbers, symbols, and whitespace.

3. Save your changes and close the editor.

4. **Optional:** Generate an upload to Hybrid Cloud Console.

```
# insights-client
```

3.4.4. Using predefined system tags to get more accurate Red Hat Insights advisor service recommendations and enhanced security

Red Hat Insights advisor service recommendations treat every system equally. However, some systems may require a higher level of security than others, or require different networking performance levels. In addition to the ability to add custom tags, Red Hat Hybrid Cloud Console provides predefined tags that the advisor service can use to create targeted recommendations for your systems that might require more attention.

To opt in and get the extended security hardening and enhanced detection and remediation capabilities offered by predefined tags, you need to configure the tags. After configuration, the advisor service provides recommendations based on tailored severity levels, and preferred network performance that apply to your systems.

To configure the tags, use the **/etc/insights-client/tags.yaml** file to tag systems with predefined tags in a similar way that you might use it to tag systems in the inventory service. To configure predefined tags, use the same **key=value** structure that you use to create custom tags. The following table lists details about the Red Hat-predefined tags.

Table 3.1. List of Supported Predefined Tags

Key	Value	Note

Key	Value	Note
security	**normal** (default) / **strict**	With **default**, the advisor service compares the system's risk profile to a baseline derived from the default configuration of the latest version of RHEL and from frequently-used usage patterns. The enables the advisor service to keep recommendations focused, actionable, and low in numbers. With the **strict**, value, the advisor service considers the system to be security-sensitive. Security-sensitive systems require specific recommendations to use a stricter baseline, potentially showing recommendations even on fresh up-to-date RHEL. installations.
network_performance`	**null** (default) / **latency** / **throughput**	The preferred network performance (either latency or throughput according to your business requirement) would affect the severity of an advisor service recommendation to a system.

NOTE

The predefined tag keys names are reserved. If you already use the key **security**, with a value that differs from one of the predefined values, you will not see a change in your recommendations. You will only see a change in recommendations if your existing **key=value** is the same as one of the predefined keys. For example, if you have a **key=value** of **security: high**, your recommendations will not change because of the Red Hat-predefined tags. If you currently have a **key=value** pair of **security: strict**, you will see a change in the recommendations for your systems.

Additional resources

- Using system tags to enable extended security hardening recommendations
- Leverage tags to make Red Hat Insights Advisor recommendations understand your environment better
- Custom system tagging

3.4.5. Configuring predefined tags

You can use the Red Hat Hybrid Cloud Console advisor service's predefined tags to adjust the behavior of recommendations for your systems. This enables your systems to gain extended security hardening and enhanced detection and remediation capabilities. This section describes how to configure the

predefined tags.

Prerequisites

- You have root-level access to your system
- You have Insights client installed
- You have systems registered within the Insights client
- You have already created the **tags.yaml** file. See Creating a tags.yaml file and adding a custom group

Procedure

- At the command line, use your preferred editor to open the **tags.yaml** configuration file located in **/etc/insights-client/**. (The following example uses Vim.)

```
[root@server ~]# vi /etc/insights-client/tags.yaml
```

- Edit the **/etc/insights-client/tags.yaml** file to add the predefined **key=value** pair for the tags. This example shows how to add **security: strict** and **network_performance: latency** tags.

```
# cat /etc/insights-client/tags.yaml
group: redhat
location: Brisbane/Australia
description:
- RHEL8
- SAP
security: strict
network_performance: latency
```

- Save your changes.
- Close the editor.
- **Optional:** Run the **insights-client** command to generate an upload to Red Hat Hybrid Cloud Console, or wait until the next scheduled Red Hat Insights upload.

```
[root@server ~]# insights-client
```

Confirming that predefined tags are in your production area

After generating an upload to Red Hat Insights (or waiting for the next scheduled Insights upload), you can check whether the tags are in the production environment by accessing Red Hat Enterprise Linux > Inventory. Find your system and look for the new tags. You should see something similar to what is shown in the following image.

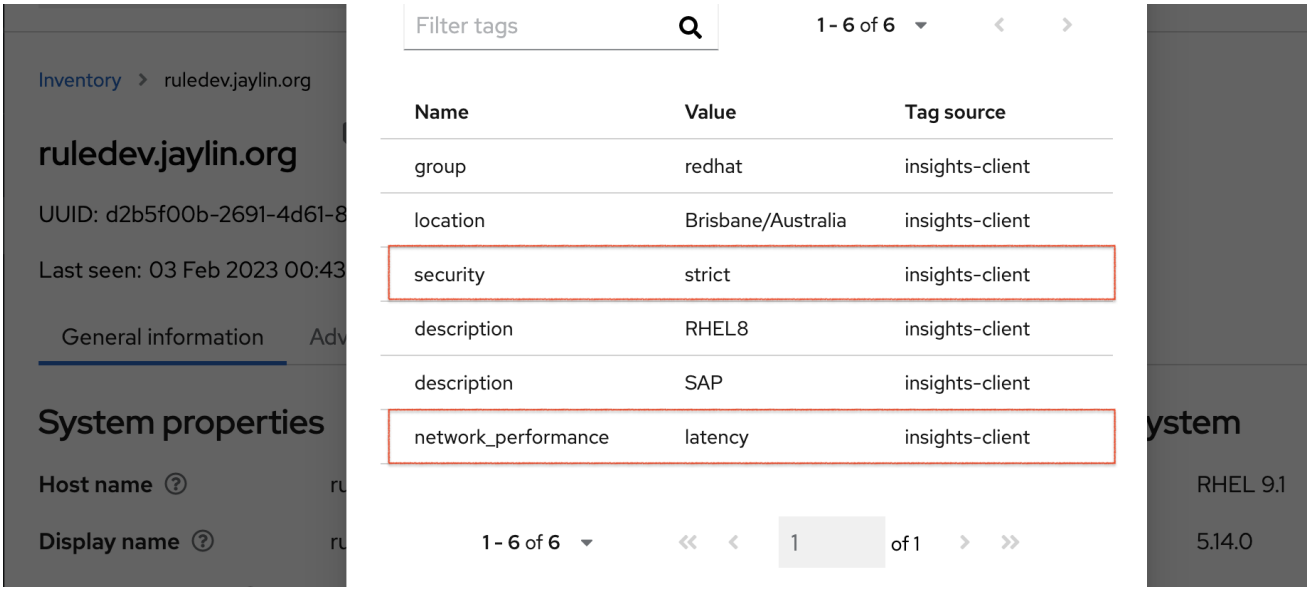

Example of recommendations after applying a predefined tag

In the following image, the advisor service shows a system with the **network_performance: latency** tag configured.

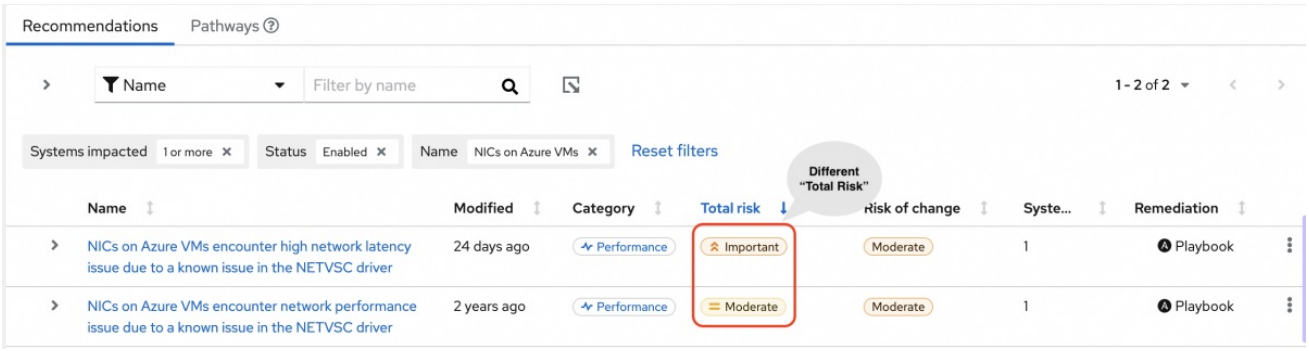

The system shows a recommendation with a higher Total Risk that is categorized as Important. The system without the **network_performance: latency** tag is categorized with a Total Risk of Moderate. You can make decisions about prioritizing the system with the higher Total Risk.

CHAPTER 4. SETTING USER PREFERENCES

Use the following procedure to set or update your email preferences.

Procedure

1. Click the user menu located on the upper-right side, then go to **User preferences → Email preferences**. The Email preferences screen opens.
   Alternatively, on the Red Hat Hybrid Cloud Console dashboard, in the left-side navigation panel at the top, click **Red Hat Enterprise Linux**, and then click **User Preferences**. The Email preferences screen opens.

2. Depending on your email notification preference, you can subscribe to **Instant notification** emails for each system with triggered policies and/or **Daily digest** (summary) of all systems with triggered policies. On this page, you can also select your preference for other Red Hat Hybrid Cloud Console emails you want to receive.

   

   NOTE

   Subscribing to instant notification can result in receiving many emails on large inventories, that is, one email per system checking in.

3. Click **Submit**.

CHAPTER 5. SYSTEM STALENESS AND DELETION

System deletion is the automated removal of systems from the Red Hat Insights inventory after all sources stop reporting information about it for a defined period of time.

System staleness is reporting when a system has missed check-ins for a defined period of time but is not yet deleted.

5.1. RULES FOR SYSTEM STALENESS AND DELETION

The inventory reporting service, as part of its messaging, includes a timestamp for when the report about the host is considered stale. This timestamp is determined by the reporting service and defaults to a value set by the user account.

When various reporters contribute data to the host in the host inventory, staleness states are recalculated.

Systems in the inventory have the following three fields related to staleness and deletion:

- **"stale_timestamp": "2019-12-13T19:36:30.979Z"**
- **"stale_warning_timestamp": "2019-12-13T19:36:30.979Z"**
- **"culled_timestamp": "2019-12-13T19:36:30.978Z"**

Rules:

- Before the stale_timestamp is reached, a system is considered fresh
- Between the stale_timestamp and stale_warning_timestamp, a system is considered stale
- Between the stale_warning_timestamp and culled_timestamp, a system is considered in the “stale warning” state and is scheduled for deletion
- After the culled_timestamp is reached, a system and all associated data are automatically deleted

5.2. VIEWING STALE SYSTEM SUMMARY

You can see stale systems and the ones scheduled for deletion in the following places in the Red Hat Insights user interface:

Dashboard: In the Red Hat Insights dashboard summary, you can see the number of stale systems and the number of systems scheduled for deletion displayed under System inventory. Click on the respective links to see the list of stale systems and the ones marked for deletion.

Inventory: On the last seen column, you will see systems marked for deletion with a warning icon. Hover on the last seen information to see if a system will be removed from the inventory in the next x days or if it is already scheduled for deletion.

5.3. FILTERING SYSTEM INVENTORY

You can filter the inventory by system status: Fresh, Stale, and Stale warning. Note that by default fresh and stale systems are listed, but stale warning systems are not shown in the user interface. You can also select the source, which is Insights by default, and filter by system name. You can remove a particular

filter in effect, or clear all at once.

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your feedback on our documentation. To provide feedback, highlight text in a document and add comments.

Prerequisites

- You are logged in to the Red Hat Customer Portal.
- In the Red Hat Customer Portal, the document is in the **Multi-page HTML** viewing format.

Procedure

To provide your feedback, perform the following steps:

1. Click the **Feedback** button in the top-right corner of the document to see existing feedback.

NOTE

The feedback feature is enabled only in the **Multi-page HTML** format.

2. Highlight the section of the document where you want to provide feedback.
3. Click the **Add Feedback** pop-up that appears near the highlighted text.
   A text box appears in the feedback section on the right side of the page.
4. Enter your feedback in the text box and click **Submit**.
   A documentation issue is created.
5. To view the issue, click the issue link in the feedback view.